The Power of Small Coalitions in Graphs

J-C. Bermond[1] D. Peleg[2]

[1] I3S CNRS, Universite de Nice-Sophia Antipolis, Bat ESSI, 930 Route des Colles, BP 145, F-06903 Sophia-Antipolis, France. E-mail: bermond@diamant.unice.fr.
[2] Department of Applied Mathematics and Computer Science, The Weizmann Institute, Rehovot 76100, Israel. E-mail: peleg@wisdom.weizmann.ac.il.

Abstract. This paper considers the question of the influence of a coalition of vertices, seeking to gain control (or majority) in local neighborhoods in a general graph. Say that a vertex v is *controlled* by the coalition M if the majority of its neighbors are from M. We ask how many vertices (as a function of $|M|$) can M control in this fashion. Upper and lower bounds are provided for this problem, as well as for cases where the majority is computed over larger neighborhoods (either neighborhoods of some fixed radius $r \geq 1$, or all neighborhoods of radii up to r). In particular, we look also at the case where the coalition must control all vertices outside itself, and derive bounds for its size.

1 Introduction

Overcoming failures is a central problem in distributed computing. A common theme in a number of approaches to this problem revolves around the notion of *majority ruling*. The idea is to eliminate the damage caused by failed vertices, or at least restrict their influence, by maintaining replicated copies of crucial data, and performing a *voting* process among the participating processors whenever faults occur, adopting the values stored at the majority of the processors as the correct data.

This method, in one form or another, is used as a component of fault-tolerant algorithms in a wide variety of contexts, including agreement and consensus problems (cf. [LSP82, Bra87, DPPU88]), quorum system applications (cf. [Gif79, GB85, OB94, JRT91]), diagnosis problems (cf. [Sul86]), self-stabilization and local mending [KP95a, KP95b], etc.

This method is usually expected to work well due to the common assumption that given today's reliable technology, at any given moment there can be only a small number of failures in the system. This implies that the required level of replication, and the extent of the voting process, can be limited.

In this paper we concentrate on understanding the majority ruling method in the context of distributed network algorithms. In this context, it is highly desirable to restrict both the replication of data stored at a processor v, and the process of majority voting regarding the data of a processor v, to processors in v's *local vicinity*.

There are two reasons for this focus on locality. First, in many cases, processors in the system are better aware of, and more involved in, whatever happens in their immediate vicinity, than far away. It is thus more natural, and often much cheaper, to store data as locally as possible. Secondly, and more importantly, the distributed network model allows only for computations which are *local* in nature, namely, in t time units, a processor can only collect data from other processors whose distance from itself in the network does not exceed t. Therefore, voting over large areas might be too expensive in terms of its time consumption.

However, there is an inherent risk in limiting ourselves to local vicinities in this way. Once replication is restricted to local neighborhoods, we run into the danger that a large enough set of faults may manage to gain the majority in some of these neighborhoods. In fact, once the voting is performed over subsets of the vertices, the ability of failed vertices to influence the outcome of the votes becomes not only a function of their *number* but also a function of their *location* in the network: well-situated vertices can acquire greater influence.

This observation naturally leads to the fundamental problem of characterizing the potential power of a set of failures in a network of processors.

Definition 1. A vertex v in a network $G(V, E)$ is said to be *controlled* by the vertex set M if the majority of its neighbors are in M.

Since our focus is on obtaining asymptotic results, there are a number of slightly different definitions for the terms "neigborhood" and "majority" that we can use in the above definition, without affecting the results. For concreteness, let us define the *neighborhood* of v as including the vertex v itself and all vertices adjacent to it, and *majority* as a strict one.

Taking the "adversarial" point of view, we formulate the following initial question:

(Q1) How many vertices (as a function of $|M|$) can a set M control?

It turns out that as far as extremal behavior is concerned, question (Q1) is easy to answer: control of virtually all vertex neighborhoods can be achieved by extremely small coalitions. As indicated by Figure 1, a set consisting of two vertices only, $M = \{a, b\}$, can gain control over the majority of the neighbors for every other vertex in $V \setminus M$.

The curious phenomenon illustrated by the above example may be viewed as an outcome of the limited scope of our majority voting. Indeed, one may hope to strengthen the quality of the voting by querying vertices to larger distances. Let $\Gamma_r(v)$ denote the r-neighborhood of v, i.e., the set of vertices at distance r or less from v (including v itself). We next pose a variant of the above question, in which neighborhoods are replaced by r-neighborhoods for some fixed r:

Definition 2. A vertex v in a network $G(V, E)$ is said to be *r-controlled* by the vertex set M if the majority of the vertices in $\Gamma_r(v)$ are in M.

(Q2) How many vertices can a set M r-control?

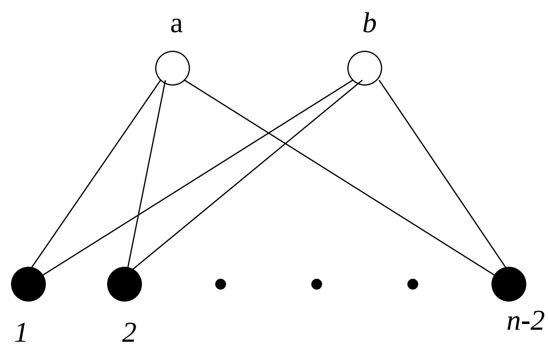

Fig. 1. A small set $M = \{a, b\}$ controlling the majority of 1-neighborhoods of all other vertices. (In all our figures, white circles represent vertices of the coalition, and black circles represent vertices controlled by the coalition. Vertices that fall under neither category, if exist, are represented by dotted circles.)

It turns out that an extremal behavior similar to that of the example of Figure 1 may occur for r-control as well, on certain graphs. More precisely, we shall present examples for every integer $r \geq 1$, in which a set M of size $r+1$ can r-control as many as $(n-r-1)/r$ vertices in an n-vertex graph G.

A more interesting picture emerges if we strengthen our voting policy, and examine *all* i-neighborhoods for a *range* of values of i.

Definition 3. A vertex v in a network $G(V, E)$ is said to be $[1, r]$-*controlled* by the vertex set M if for every $1 \leq i \leq r$, the majority of the vertices in $\Gamma_i(v)$ are in M.

(Q3) How many vertices can a set M $[1, r]$-control?

Our results imply that the answer to this last question is $O(|M|^{1+1/\lfloor \log_2 r \rfloor})$, and that this result is tight, in the sense that there exist (infinitely many) graphs and sets M that achieve this influence.

A special case of the above problems was raised and studied in [LPRS93]. It is based on the following notion.

Definition 4. Call the set M an r-*monopoly* (respectively $[1, r]$-*monopoly*) if it r-controls (respectively $[1, r]$-*controls*) every vertex in the graph.

The question addressed in [LPRS93] was:

What can be said about the size of monopolies in the graph?

Tight answers to this question was provided in [LPRS93], by relating it to some natural packing and covering problems in graphs. Specifically, the following results were established in [LPRS93].

Proposition 5. **[LPRS93]**

1. *In every n-vertex graph, a 1-monopoly must be of size* $\Omega(\sqrt{n})$.
2. *There exist (infinitely many) n-vertex graphs with 1-monopolies of size* $O(\sqrt{n})$.

A graph G^α with a 1-monopoly of size $O(\sqrt{n})$ as in Prop. 5(2) is depicted in Fig. 2. The graph consists of a coalition M of $2\sqrt{n}$ vertices, u_i, w_i for $1 \le i \le \sqrt{n}$, connected by a clique. The rest of the vertices are partitioned into $\sqrt{n}$ groups of $\sqrt{n}-2$ vertices each, where the vertices of the i'th group are attached to u_i and w_i.

As for r-monopolies, it is shown in [LPRS93] that for any fixed $r \ge 1$ there exist (infinitely many) n-vertex graphs with r-monopolies of size $O(n^{2/3})$. This bound on the size of r-monopolies was very recently improved [BBPP95], and the new bound is tight for even r.

Proposition 6. **[BBPP95]**

1. *In every n-vertex graph, an r-monopoly for even* $r \ge 2$ *must be of size* $\Omega(n^{3/5})$.
2. *For any fixed* $r \ge 2$ *there exist (infinitely many) n-vertex graphs with r-monopolies of size* $O(n^{3/5})$.

Proposition 7. **[LPRS93]**

1. *For every fixed* $r \ge 1$ *and n-vertex graph, any* $[1,r]$*-monopoly must be of size* $\Omega(n^{1-1/(\lfloor \log_2 r \rfloor + 2)})$.
2. *For every fixed* $r \ge 1$ *there exist (infinitely many) n-vertex graphs with* $[1,r]$*-monopolies of size* $O(n^{1-1/(\lfloor \log_2 r \rfloor + 2)})$.

Our bounds on the extent of control possible for a set of vertices are derived using a modified variant of the integral packing technique developed in [LPRS93].

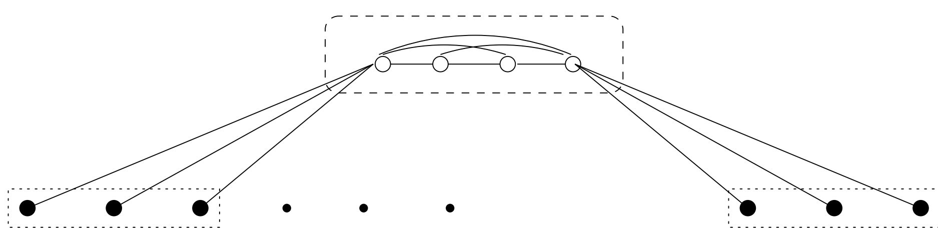

Fig. 2. The graph G^α with a 1-monopoly M of size $O(\sqrt{n})$. Here each white circle represents a pair of vertices u_i and w_i.

While the surprising power of small coalitions is clearly demonstrated in the results of [LPRS93], there are settings in which controlling coalitions can be even smaller. In particular, in a context where we think of the coalition seeking control as a set of faulty (possibly malicious) processors, it may as well be assumed that the coalition M is only interested in gaining control over the neighborhoods of *other* vertices, belonging to $V \setminus M$. This is because the vertices in the coalition are not obligated by the rules of the "voting game" anyhow, so the "adversary" needs not "waste its powers" (so to speak) on controlling them. Such a coalition can therefore be considerably smaller. For example, the set $\{a, b\}$ in Figure 1, controls every vertex in $V \setminus M$, in sharp contrast with Prop. 5(1).

More generally, we can define the following notion.

Definition 8. A *self-ignoring r-monopoly* M is a set that r-controls every vertex in $V \setminus M$ (and similarly for a self-ignoring $[1, r]$-monopoly).

We can now repeat the questions of [LPRS93] for self-ignoring monopolies. It turns out that the results have a rather similar structure, except "shifted" downwards. In particular, for a self-ignoring r-monopoly in an n-vertex graph, we prove a lower bound of $|M| = \Omega(n^{1/2})$ for fixed $r \geq 2$. (For $r = 1$, the example of Fig. 1 prevents any such lower bound.) This lower bound is complemented by examples establishing the existence of n-vertex graphs with self-ignoring r-monopolies of size $O(n^{1/2})$ for every fixed $r \geq 2$.

Turning to self-ignoring $[1, r]$-monopolies, it follows from our bounds on the extent of control possible for vertex sets, that any self-ignoring $[1, r]$-monopoly for an n-vertex graph G must contain at least $\Omega(n^{1-1/(\lfloor \log_2 r \rfloor + 1)})$ vertices, and that there are n-vertex graphs with self-ignoring $[1, r]$-monopolies of size $O(n^{1-1/(\lfloor \log_2 r \rfloor + 1)})$, so this bound is tight.

Note again the slight difference in the exponent between the bounds for $[1, r]$-monopolies and self-ignoring $[1, r]$-monopolies (see Fig. 3).

A bound of the form $\Omega(n^{1-1/(\lfloor \log_2 r \rfloor + 1)})$ can be derived for both cases using (slightly different variants of) the integral packing technique developed in [LPRS93]. This bound is tight for the self-ignoring case, but for the case of a full monopoly, proving Prop. 7(1) is done via a different technique for constructing fractional packings.

max. radius r	$[1, r]$-monopoly	self-ignoring $[1, r]$-monopoly
1	$\Theta(n^{1/2})$	$\Theta(1)$
2,3	$\Theta(n^{2/3})$	$\Theta(n^{1/2})$
2^{t-1} to $2^t - 1$	$\Theta(n^{1-1/(t+1)})$	$\Theta(n^{1-1/t})$

Fig. 3. Size comparison of $[1, r]$-monopolies vs. self-ignoring $[1, r]$-monopolies.

Certain *dynamic* variants of majority voting problems were studied in the literature, in the context of discrete time dynamical systems. These variants concentrated on a setting in which the nodes of the system operate in discrete time steps, and at each step, each node computes the majority in its neighborhood, and adapts the resulting value as its own. The typical problems studied in this setting involve the behavior of the resulting sequence of global states (represented as a vector $\bar{x}^t = (x_1^t, \ldots, x_n^t)$, where x_i^t represents the value at node v_i after time step t). For instance, the fact that the period of such sequences is either one or two is proved (in various contexts) in [GO80, PS83, PT86]. The problem was studied further in [Mor94c, Mor94b, Mor94a]. Also, the applicability of majority voting as a tool for fault-local mending was investigated in [KP95a, KP95b].

2 $[1, r]$-controlling coalitions

Definition 9. Given an n-vertex graph G, a *packing* is a collection $\mathcal{P} = \{P_1, \ldots, P_t\}$ of disjoint neighborhoods in G. For each neighborhood P_i, we denote its center by c_i and its radius by r_i (namely, $P_i = \Gamma_{r_i}(c_i)$ for every i). The *volume* of $\mathcal{P}$ is defined as $\mathcal{V}(\mathcal{P}) = \sum_i |P_i|$.

Definition 10. Given a set of vertices X, a packing $\mathcal{P}$ is said to be X-*centered* if all the centers of its neighborhoods are from X.

We make use of the following lemma, which is an extension of a lemma of [LPRS93].

Lemma 11. *For every n-vertex graph G, set of vertices X and fixed integer r, there exists an X-centered packing $\mathcal{P}$ in G, with neighborhoods of radius at most r, and volume $\mathcal{V}(\mathcal{P}) \geq |X|^{1-1/(\lfloor \log_2 r \rfloor + 1)}$. All neighborhoods in the packing may be restricted to have a radius which is a power of 2.*

Proof: Let $t = \lfloor \log_2 r \rfloor$ and $q = 1/(t+1)$, and for every $1 \leq i \leq t+1$ let $k_i = 2^{t-i+1}$. We now construct an X-centered packing $\mathcal{P}$ for G as follows.

As an initial test, if there is some vertex v in X such that $|\Gamma_{k_1}(v)| \geq |X|^{1-q}$, then we are done, since by assumption, this neighborhood alone contains $|X|^{1-q}$ vertices, and its radius is $k_1 \leq r$, hence it can serve as the desired X-centered packing $\mathcal{P}$. Thus for the sequel, assume that $|\Gamma_{k_1}(v)| < |X|^{1-q}$ for every $v \in X$.

The packing proceeds in an iterative process consisting of t phases, $i = 1, \ldots, t$. Intuitively, each phase i adds to $\mathcal{P}$ a collection of disjoint k_{i+1}-neighborhoods, and eliminates vertices in their vicinity from consideration for later iterations. In particular, throughout the execution we maintain a set S containing those vertices of X whose neighborhoods have not yet been considered. The process relies on the property that the following precondition holds at the beginning of each phase.

(PRE$_i$) At the beginning of phase i, $|\Gamma_{k_i}(v)| < |X|^{1-iq}$ for every vertex v in S.

For $i = 1$, S is set to be X, and the assumed condition (PRE$_1$) holds as discussed above.

Phase i of the packing process operates as follows. As long as S contains vertices v such that $|\Gamma_{k_{i+1}}(v)| \geq |X|^{1-(i+1)q}$, pick one such vertex v arbitrarily, add $\Gamma_{k_{i+1}}(v)$ to the constructed packing $\mathcal{P}$, and eliminate from S every vertex v' whose neighborhood $\Gamma_{k_{i+1}}(v')$ intersects $\Gamma_{k_{i+1}}(v)$ (including v itself). Note that the condition for eliminating v' may also be stated as $dist(v, v') \leq 2k_{i+1} = k_i$.

We now make the following observations. First, observe that the assumption (PRE$_i$) on the sizes of neighborhoods at the beginning of each phase i is guaranteed by the halting condition of the previous phase. Secondly, note that after t phases, the vertices of the set X will all be exhausted. Also observe that the neighborhoods added to $\mathcal{P}$ are all disjoint. This holds because once a vertex w is included in some neighborhood $\Gamma_{k_{i+1}}(v)$ at some phase i, we eliminate from S every vertex v' whose k_{i+1}-neighborhood contains w. Of course w still appears in *larger* neighborhoods of other vertices remaining in S, but later phases consider only neighborhoods of *equal or smaller* radius.

In order to analyze the volume of the resulting packing $\mathcal{P}$, let us consider some vertex v_j whose neighborhood is picked to $\mathcal{P}$ at some phase i, and denote the cardinality of its k_{i+1}-neighborhood by m_j, and the number of vertices v' eliminated because of it from S by l_j. But as argued earlier, each such v' was eliminated because $dist(v, v') \leq k_i$, or $v' \in \Gamma_{k_i}(v_j)$. This implies that

$$l_j \;\leq\; |\Gamma_{k_i}(v_j)| \;<\; |X|^{1-iq}. \tag{1}$$

(The latter inequality follows from the fact that v_j has not been selected in the $(i-1)$'st phase, hence it obeys condition (PRE$_i$).) It is also clear that

$$m_j \;=\; |\Gamma_{k_{i+1}}(v_j)| \;\geq\; |X|^{1-(i+1)q}. \tag{2}$$

(The inequality follows from the fact that v_j *has* been selected in the i'th phase.) It follows from inequalities (1) and (2) that $l_j < |X|^q \cdot m_j$, and therefore the total volume satisfies

$$\mathcal{V}(\mathcal{P}) \;=\; \sum_j m_j \;\geq\; \frac{1}{|X|^q}\sum_j l_j \;=\; |X|/|X|^q,$$

completing the proof of the lemma. ∎

We now derive a bound on the maximum number of vertices that can be controlled by a coalition M.

Theorem 12. *In every graph G and for every fixed integer $r \geq 2$, a coalition M can $[1, r]$-control at most $O(|M|^{1+1/\lfloor \log_2 r \rfloor})$ vertices.*

Proof: Let X be the set of vertices that are $[1, r]$-controlled by the coalition M. By Lemma 11, there exists an X-centered packing $\mathcal{P}$ in G, with neighborhoods of radius at most r, and volume $\mathcal{V}(\mathcal{P}) \geq |X|^{1-1/(\lfloor \log_2 r \rfloor + 1)}$. Since the vertices of X are $[1, r]$-controlled by M, each of the neighborhoods in $\mathcal{P}$ contains a majority of vertices from M. By the fact that the neighborhoods in the packing $\mathcal{P}$ are disjoint, $|M| > \frac{1}{2}|X|^{1-1/(\lfloor \log_2 r \rfloor + 1)}$. The claim follows. ∎

Theorem 12 implies that the number of vertices that can be $[1,r]$-controlled by the coalition M for $r=2$ or $r=3$ is at most $|M|^2$. For $4 \le r \le 7$, that number is bounded by $|M|^{1.5}$, etc.

Corollary 13. *In every graph G and for every fixed integer $r \ge 1$, any self-ignoring $[1,r]$-monopoly M must be of cardinality $|M| = \Omega(n^{1-1/(\lfloor \log_2 r \rfloor + 1)})$.*

Proof: For $r=1$ the claim holds trivially. For $r \ge 2$, M must satisfy $V \setminus M = O(|M|^{1+1/\lfloor \log_2 r \rfloor})$ by Theorem 12, and the claim follows. ∎

The bounds of Theorem 12 and Corollary 13 are tight. In fact, the proof for the existence of a small self-ignoring $[1,r]$-monopoly M is based on the same example graph $G_{t,p}$ constructed in [LPRS93] for establishing the upper bound of Prop. 7(2). The required case analysis is slightly different, though, given that the bound proved is different too. This analysis is postponed to the full paper.

Theorem 14. *For every fixed integer $r \ge 1$ there exist (infinitely many) n-vertex graphs G_n and self-ignoring $[1,r]$-monopolies M_n in G_n, such that $|M_n| = \Theta(n^{1-1/(\lfloor \log_2 r \rfloor + 1)})$.*

Theorem 14 establishes the tightness of the bound of Corollary 13. As a straightforward corollary, we get that the bound of Theorem 12 is tight as well.

Corollary 15. *For every fixed integer $r \ge 1$ there exist (infinitely many) n-vertex graphs G_n and coalitions M_n, such that M_n $[1,r]$-controls $\Theta(|M_n|^{1+1/\lfloor \log_2 r \rfloor})$ vertices in G_n.* ∎

3 r-controlling coalitions

For r-control, the situation is different, in that a very small set M can r-control a very large set of vertices. For $r=1$ this is demonstrated by the example in Fig. 1, where a set of size $|M|=2$ achieves 1-control over the remaining $n-2$ vertices. This example can be generalized to show the following.

Theorem 16. *For any integer r there exists a family of n-vertex graphs G_n and sets M_n, such that M_n r-controls a subset X_n of $V \setminus M_n$, and $|M_n| = r+1$, $|X_n| = (n-r-1)/r$.*

Proof: Given r and p, let $n = rp + (r+1)$. Construct $G^{\gamma}_{r,p}$ as follows. The graph is *leveled*, namely, the vertices are arranged into $r+1$ levels, numbered 1 through $r+1$, with edges connecting only vertices in adjacent levels ℓ, $\ell+1$. Each level $2 \le \ell \le r+1$ contains p vertices, $v^{\ell}_1, \ldots, v^{\ell}_p$, and level 1 contains $r+1$ vertices. Let X denote the set of vertices on level $r+1$, and let M denote the set of vertices on level 1. When p is very large with respect to r, X contains roughly a $1/r$ fraction of the vertices of the graph, yet the edge connections defined next will guarantee that M has the majority in any r-neighborhood around the vertices of X.

The edges connecting two consecutive levels $\ell-1$ and ℓ are defined as follows. The vertices of level 1 (M) are connected by a complete bipartite graph (cross-bar) to the vertices of level 2. From level 2 and on, the vertices of the different levels form chains of length r. Namely, for $2 \le \ell \le r$, each vertex v_i^ℓ of level ℓ is connected to vertex $v_i^{\ell+1}$ of level $\ell+1$. Figure 4 depicts an example graph $G^\gamma_{r,p}$ for $r=3$ and some p.

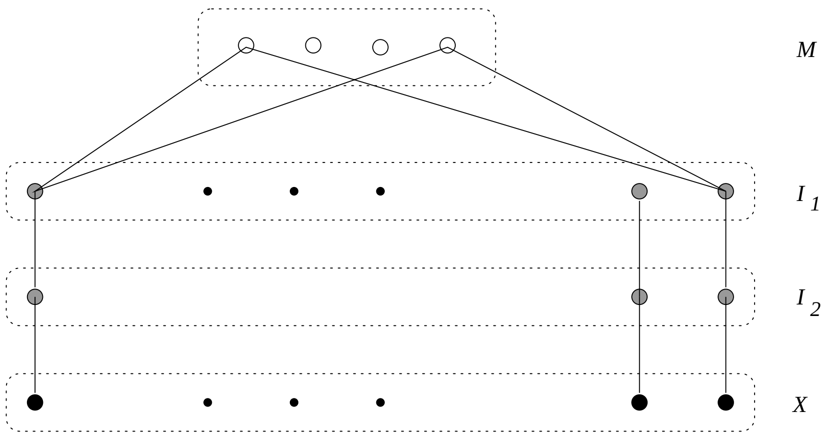

Fig. 4. The graph $G^\gamma_{r,p}$ for $r = 3$ and some p, with a set M of size 4 controlling the majority of 3-neighborhoods of the vertices of a set X of size $p = (n-4)/3$.

It is straightforward to verify that the vertices of M r-control those of X. ∎

Finally we turn to self-ignoring r-monopolies. As mentioned in the introduction, there exist graphs with self-ignoring 1-monopolies of size 2 (see Fig. 1). We now derive a lower bound on the size of self-ignoring r-monopolies for $r \ge 2$. For the proof we use the following notation. Let $S_i = \bigcup_{v\in M} \Gamma_i(v)$ for every $i \ge 1$.

Theorem 17. *In every graph G and for every fixed $r \ge 2$, any self-ignoring r-monopoly M must be of cardinality $|M| = \Omega(n^{1/2})$.*

Proof: Consider a graph G and a self-ignoring r-monopoly M, for $r \ge 2$. Let

$$S_i = \bigcup_{v\in M} \Gamma_i(v), \qquad 1 \le i \le r.$$

Let $M_1 = S_1 \setminus M$ and $M_i = S_i \setminus S_{i-1}$. Any vertex of V is at distance at most r from M, so $V \setminus M = \bigcup_{1\le i\le r} M_i$.

Fact 18 *For every $1 \le i \le r-1$ and $v \in M$, $|\Gamma_i(v) \cap M_i| < |M|$.*

Proof: Let v be a vertex of M. If v has no neighbors in M_1, then $\Gamma_i(v) \cap M_i = \emptyset$, and we are done. So now suppose v has such neighbors, and let w be a vertex in M_1 adjacent to v. If $|\Gamma_i(v) \cap M_i| \geq |M|$, then the r-neighborhood $\Gamma_r(w)$ (which clearly contains all of $\Gamma_i(v) \cap M_i$) would contain more than $|M|$ vertices, contradicting the fact that this neighborhood should contain a majority of vertices from M. ∎

From Fact 18 we deduce:

Fact 19 *(a) For every $1 \leq i \leq r-1$, $|M_i| < |M|^2$.*
(b) For every $1 \leq i \leq r-1$, the number of pairs (v, w) with $v \in M$ and $w \in M_i \cap \Gamma_i(v)$ is at most $|M|^2$.

Now let us prove that $|M_r| < |M|^2$. We will use the following relations. For every vertex $u_j \in M_r$ let $A_j = \Gamma(u_j) \cap M_{r-1}$ and $B_j = \bigcup_{w \in A_j} \Gamma(w) \cap M_r$, and let $C(M, A_j)$ denote the set of pairs (v, w) such that $v \in M$ and $w \in A_j \cap \Gamma_{r-1}(v)$ (or by symmetry $v \in \Gamma_{r-1}(w)$). We have $\Gamma_r(u_j) \cap M = \{v \in M \mid \text{ there exists a } w \in A_j \text{ such that } (v, w) \in C(M, A_j)\}$. As each $v \in \Gamma_r(u_j \cap M)$ is the first element of at least one pair of $C(M, A_j)$, $|\Gamma_r(u_j) \cap M| \leq |C(M, A_j)|$. Since $r \geq 2$, $B_j \subseteq \Gamma_r(u_j) \cap M_r$. Therefore, as the r-neighborhood of u_j should contain a majority of vertices of M,

$$|B_j| \leq |\Gamma_r(u_j) \cap M| \leq |C(M, A_j)|. \tag{3}$$

Now let us select a sequence of vertices from M_r, denoted $u_1, u_2, \ldots, u_p$, as follows. First, pick u_1 to be an arbitrary vertex in M_r. If $B_1 \neq M_r$, then pick u_2 to be an arbitrary vertex in $M_r \setminus B_1$. Repeat this process as long as $\bigcup_{1 \leq k \leq j} B_k \neq M_r$, picking u_{j+1} to be an arbitrary vertex in $M_r \setminus \bigcup_{i \leq k \leq j} B_k$.

At the end of this process, we have

$$\bigcup_{1 \leq j \leq p} B_j = M_r. \tag{4}$$

By the definition of the B_j's, all the A_j's associated with the chosen u_j are pairwise disjoint. Indeed, $u_j \notin \bigcup_{1 \leq k \leq j-1} B_k$, and so it cannot be adjacent to any vertex of the A_k, $1 \leq k \leq j-1$.

As the A_j are pairwise disjoint, $C(M, A_i) \cap C(M, A_j) = \emptyset$ for $i \neq j$. So we have

$$\sum_j |C(M, A_j)| \leq |C(M, M_{r-1})|. \tag{5}$$

Combining equation (4), inequalities (3) and (5) and Fact 19(b), we conclude

$$|M_r| = |\bigcup_{1 \leq j \leq p} B_j| \leq \sum_{1 \leq j \leq r} |B_j| \leq \sum_{1 \leq j \leq r} |C(MA_j)| \leq C(M_1, M_{r-1}) \leq |M|^2.$$

In summary, $|M_i| < |M|^2$ for $1 \leq i \leq r$, and hence $|V - M| \leq r|M|^2$, and the theorem follows. ∎

As can be seen from the following theorem, this bound is tight.

Theorem 20. *For every fixed integer $r \geq 2$ there exist (infinitely many) n-vertex graphs G_n and self-ignoring r-monopolies M_n in G_n, such that $|M_n| = \Theta(n^{1/2})$.*

Proof: For $r = 2$, we note that in the graph G^{α} of Fig. 2, the coalition M (presented there as a 1-monopoly) is also a self-ignoring 2-monopoly.

Next, we give an example of such a set for any $r > 2$. For integers r, p, construct the graph $G^{\delta}_{r,p}$ as follows. The graph is again *leveled*, namely, the vertices are arranged into $\lfloor r/2 \rfloor + 2$ levels, numbered 1 through $\lfloor r/2 \rfloor + 2$, with edges connecting only vertices in adjacent levels ℓ, $\ell + 1$. Level 1 contains p^2 vertices, each level $2 \leq \ell \leq \lfloor r/2 \rfloor + 1$ contains p vertices, and level $\lfloor r/2 \rfloor + 2$ contains a single vertex. Let X denote the set of vertices on level 1, and let $M = V \setminus X$ be the rest of the vertices. When p is much larger than r, M contains roughly $\sqrt{n}$ of the vertices of the graph, yet our construction will have the property that the vertices of M majorize all r-neighborhoods of X vertices.

The edges connecting two consecutive levels $\ell - 1$ and ℓ are now defined as follows. The single vertex of level $\lfloor r/2 \rfloor + 2$ is connected to all the vertices of level $\lfloor r/2 \rfloor + 1$. For level $2 \leq \ell \leq \lfloor r/2 \rfloor + 1$, each vertex v is connected to the corresponding vertex at level $\ell - 1$. For level $\ell = 2$, each vertex of level 2 has p distinct neighbors at level 1 (i.e., each vertex of X has exactly one neighbor on level 2). See Figure 5 for an example graph $G^{\delta}_{r,p}$ for $r = 5$, $p = 5$.

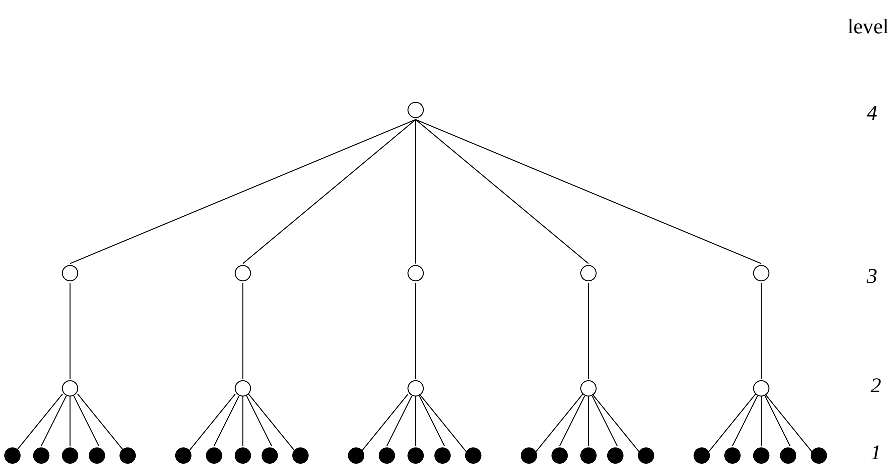

Fig. 5. The graph $G^{\delta}_{r,p}$ for $r = 5$ and $p = 5$.

A straightforward case analysis reveals that in the graph $G^{\delta}_{r,p}$, for every vertex $v \in X$, the majority of the vertices in $\Gamma_r(v)$ are from M. $\blacksquare$

References

[BBPP95] J-C. Bermond, J. Bond, D. Peleg, and S. Perennes. Tight bounds on the size of 2-monopolies. In preparation, 1995.

[Bra87] G. Bracha. An $o(\log n)$ expected rounds randomized Byzantine generals algorithm. *J. of the ACM*, pages 910–920, 1987.

[DPPU88] C. Dwork, D. Peleg, N. Pippenger, and E. Upfal. Fault tolerance in networks of bounded degree. *SIAM J. on Comput.*, 17:975–988, 1988.

[GB85] H. Garcia-Molina and D. Barbara. How to assign votes in a distributed system. *J. ACM*, 32(4):841–860, 1985.

[Gif79] D. K. Gifford. Weighted voting for replicated data. In *Proc. 7th Symp. Oper. Sys. Princip.*, pages 150–159, 1979.

[GO80] E. Goles and J. Olivos. Periodic behaviour of generalized threshold functions. *Discrete Mathematics*, 30:187–189, 1980.

[JRT91] P. Jalote, S. Rangarajan, and S. K. Tripathi. Capacity of voting systems. Technical Report UMIACS-TR-91-118, University of Maryland, 1991.

[KP95a] Shay Kutten and David Peleg. Fault-local distributed mending. In *Proc. 14th ACM Symp. on Principles of Distributed Computing*, August 1995.

[KP95b] Shay Kutten and David Peleg. Tight fault-locality. In *Proc. 36th IEEE Symp. on Foundations of Computer Science*, October 1995.

[LPRS93] N. Linial, D. Peleg, Y. Rabinovich, and M. Saks. Sphere packing and local majorities in graphs. In *2nd ISTCS*, pages 141–149. IEEE Computer Soc. Press, June 1993.

[LSP82] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, July 1982.

[Mor94a] G. Moran. On the period-two-property of the majority operator in infinite graphs. Manuscript, 1994.

[Mor94b] G. Moran. Parametrization for stationary patterns of the r-majority operators on 0-1 sequences. *Discrete Mathematics*, 132:175–195, 1994.

[Mor94c] G. Moran. The r-majority vote action on 0-1 sequences. *Discrete Mathematics*, 132:145–174, 1994.

[OB94] M. Obradovic and P. Berman. Voting as the optimal static pessimistic scheme for managing replicated data. *IEEE Trans. on Parallel & Distr. Systems*, 5(1):64–73, 1994.

[PS83] Svatopluk Poljak and Miroslav Sura. On periodical behaviour in societies with symmetric influences. *Combinatorica*, 3:119–121, 1983.

[PT86] Svatopluk Poljak and Daniel Turzik. On an application of convexity to discrete systems. *Discrete Applied Mathematics*, 13:27–32, 1986.

[Sul86] G.F. Sullivan. *The complexity of system-level fault diagnosis and diagnosability.* PhD thesis, Department of Computer Science, Yale University, 1986.

This article was processed using the LaTeX macro package with SIROCCO style